



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/769,388	01/30/2004	Colyn S. Case	19680-007200US	9441

20350 7590 04/19/2006

TOWNSEND AND TOWNSEND AND CREW, LLP
TWO EMBARCADERO CENTER
EIGHTH FLOOR
SAN FRANCISCO, CA 94111-3834

EXAMINER

GOLDEN, JAMES R

ART UNIT PAPER NUMBER

2187

DATE MAILED: 04/19/2006

Please find below and/or attached an Office communication concerning this application or proceeding.

Office Action Summary	Application No. 10/769,388	Applicant(s) CASE ET AL.	
	Examiner James Golden	Art Unit 2187	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) ☒ Responsive to communication(s) filed on 30 January 2004.
- 2a) ☐ This action is **FINAL**. 2b) ☒ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) ☒ Claim(s) 1-26 is/are pending in the application.
- 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) ☐ Claim(s) _____ is/are allowed.
- 6) ☒ Claim(s) 1-26 is/are rejected.
- 7) ☐ Claim(s) _____ is/are objected to.
- 8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) ☐ The specification is objected to by the Examiner.
- 10) ☒ The drawing(s) filed on 30 January 2004 is/are: a) ☐ accepted or b) ☒ objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) ☐ All b) ☐ Some * c) ☐ None of:
1. ☐ Certified copies of the priority documents have been received.
2. ☐ Certified copies of the priority documents have been received in Application No. _____.
3. ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).
- * See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- | | |
|---|---|
| 1) <input checked="" type="checkbox"/> Notice of References Cited (PTO-892) | 4) <input type="checkbox"/> Interview Summary (PTO-413)
Paper No(s)/Mail Date. _____ |
| 2) <input type="checkbox"/> Notice of Draftsperson's Patent Drawing Review (PTO-948) | 5) <input type="checkbox"/> Notice of Informal Patent Application (PTO-152) |
| 3) <input checked="" type="checkbox"/> Information Disclosure Statement(s) (PTO-1449 or PTO/SB/08)
Paper No(s)/Mail Date <u>07 June 2004</u> . | 6) <input type="checkbox"/> Other: _____ |

DETAILED ACTION

The instant application 10/769388 has a total of 26 claims pending. There are 2 independent claims and 24 dependent claims.

Information Disclosure Statement

1. The information disclosure statement submitted on 06/07/2004 is in compliance with the provisions of 37 CFR 1.97. Accordingly, the information disclosure statement is being considered by the examiner.

Drawings

2. The drawings are objected to as failing to comply with 37 CFR 1.84(p)(5) because they do not include the following reference sign mentioned in the description: 306 of Fig. 3 (page 15, line 19).

3. This may not be an exhaustive list of mistakes in the drawings, and the examiner respectfully requests that the applicant check the drawings for other mistakes.

4. Corrected drawing sheets in compliance with 37 CFR 1.121(d) are required in reply to the Office action to avoid abandonment of the application. Any amended replacement drawing sheet should include all of the figures appearing on the immediate prior version of the sheet, even if only one figure is being amended. The figure or figure number of an amended drawing should not be labeled as "amended." If a drawing figure is to be canceled, the appropriate figure must be removed from the replacement sheet, and where necessary, the remaining figures must be renumbered and appropriate

changes made to the brief description of the several views of the drawings for consistency. Additional replacement sheets may be necessary to show the renumbering of the remaining figures. Each drawing sheet submitted after the filing date of an application must be labeled in the top margin as either "Replacement Sheet" or "New Sheet" pursuant to 37 CFR 1.121(d). If the changes are not accepted by the examiner, the applicant will be notified and informed of any required corrective action in the next Office action. The objection to the drawings will not be held in abeyance.

Claim Objections

5. **Claims 1-21** are objected to because of the following informalities: Claim 1 recites the limitation "the input arbitration circuit" (line 10). There is insufficient antecedent basis for this limitation in the claim. Claims 2-21 are objected to because of their dependence on claim 1. These objections can be overcome by correcting line 3 of claim 1 to --an input arbitration circuit-- or by correcting all instances of "input arbitration circuit" to read --input arbitration module--.

b

Claim Rejections - 35 USC § 103

6. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

7. **Claims 1-2, 5-7, 10-12, 22 and 24-25** are rejected under 35 U.S.C. 103(a) as being unpatentable over Kikuta et al. (US 6,260,131) in view of Carter et al. (US 6,003,123).

8. **With respect to claim 1**, Kikuta et al. disclose a translation lookaside buffer system for translating virtual addresses to physical addresses, the system comprising:

- an input arbitration module (Memory Management Unit [MMU] 122 of Fig. 2) configured to receive translation requests from a plurality of clients (Figs. 7 and 9 show a plurality of clients as processors 1-3) each translation request including a target virtual address, the input arbitration module further configured to select one of the translation requests for processing (column 9, lines 28-33, the MMU “intercepts” a memory access request, thereby selecting it);
- a primary cluster store (TLB 126 of Figs. 2 and 4; global TLB of Figs. 7 and 9) having a plurality of locations, each location configured to store a cluster, wherein each cluster provides a mapping to a respective physical address for each virtual address in a range of virtual address space (column 3, lines 36-43; column 8, line 66 -- column 9, line 6; a TLB entry represents a location, and each entry maps a virtual address to its corresponding physical address where the virtual address space is represented by the virtual page number);
- a primary lookup logic circuit (MMU 122 of Fig. 2) coupled to receive the selected translation request from the input arbitration circuit (column 9, lines 28-33) and configured to associate one of the locations in the primary cluster store with the selected translation request (column 10, lines 4-5), thereby designating the

associated location for storing a cluster whose range includes the target virtual address (column 10, lines 4-5); and

- a translation module (MMU 122 of Fig. 2) configured to translate the target virtual address of the selected translation request to a physical address based at least in part on the mapping provided by the cluster stored in the primary cluster store location associated with the selected translation request (column 9, lines 33-35).

Kikuta et al. do not disclose the limitation wherein the respective ranges of different ones of the clusters have different sizes.

However, Carter et al. disclose the limitation wherein the respective ranges of different ones of the clusters have different sizes (column 4, lines 40-46; column 6, lines 32-35).

Kikuta et al. and Carter et al. are analogous art because they are from the same field of endeavor, namely virtual-to-physical address translation in a multiprocessor system.

At the time of invention, it would have been obvious to a person of ordinary skill in the art to combine the TLB system of Kikuta et al. with the variable segment size of Carter et al. The motivation for doing so would have been because "this segmentation and access control system provides flexibility to the user, while still permitting strictly enforced security" (column 6, lines 42-44).

Therefore, it would have been obvious to a person of ordinary skill in the art to combine Kikuta et al. with Carter et al. for the benefit of a TLB system in a network with variably-sized segments to obtain the invention as specified in claim 1.

9. **With respect to claim 2**, Kikuta et al. in view of Carter et al. disclose the translation lookaside buffer system of claim 1 (see above paragraph 8). Kikuta et al. disclose the limitations further comprising:

- a primary tag store (TLB 126 of Figs. 2 and 4; global TLB of Figs. 7 and 9) coupled to the primary lookup logic circuit (TLB 126 of Fig. 2 is onboard MMU 122 of Fig. 2 and therefore coupled to it) and configured to store a tag corresponding to each location in the primary cluster store (column 3, lines 36-43; column 8, line 66 -- column 9, line 6, where the tag is the virtual page number, and each location in the TLB is supplied with an entry 140 of Fig. 3), wherein each tag encodes the range of the virtual address space mapped by the cluster in the corresponding primary cluster store location (column 3, lines 15-18),
- wherein the primary lookup logic circuit (MMU 122 of Fig. 2) is further configured to detect a primary cluster hit in the event that the target virtual address of the selected translation request matches one of the tags in the primary tag store (column 9, line 33) and to detect a primary cache miss otherwise (column 9, lines 47-50).

10. **With respect to claim 5**, Kikuta et al. in view of Carter et al. disclose the translation lookaside buffer system of claim 2 (see above paragraph 9). Kikuta et al. disclose the limitations wherein the primary lookup logic circuit (MMU 122 of Fig. 2) is further configured to respond to a primary cache hit by associating the location in the

primary cluster store that corresponds to the matched tag with the selected translation request (column 9, lines 33-38).

11. **With respect to claim 6**, Kikuta et al. in view of Carter et al. disclose the translation lookaside buffer system of claim 5 (see above paragraph 10). Kikuta et al. disclose the limitations further comprising a secondary cluster cache (page table 129 of Fig. 2, column 9, line 58) configured to respond to a detected primary cache miss by obtaining a cluster based on the target virtual address (column 9, lines 57-66) and storing the obtained cluster in the primary cluster store (column 9, line 66 -- column 10, line 5).

12. **With respect to claim 7**, Kikuta et al. in view of Carter et al. disclose the translation lookaside buffer system of claim 6 (see above paragraph 11). Kikuta et al. disclose the limitations wherein the secondary cluster cache includes:

- a secondary cluster store configured to store a plurality of clusters (page table 129 of Fig. 2);
- a secondary tag store (internal to page table 129 of Fig. 2) configured to store a secondary tag corresponding to each cluster stored in the secondary cluster store (column 3, lines 36-43; column 8, line 66 -- column 9, line 6, lines 57-66, where the tag is the virtual page number), wherein each secondary tag encodes a range of virtual addresses mapped by the corresponding cluster (column 3, lines 15-18); and
- a secondary lookup logic circuit (PMH 124 of Fig. 2; column 9, lines 47-50) configured to receive the target virtual address (column 9, lines 51-53) and to

identify a matching one of the secondary tags in the secondary tag store (column 9, lines 57-59), wherein the range of virtual addresses encoded by the matching secondary tag includes the current target virtual address (column 9, lines 54-55; the range encoded by the tag is defined by the virtual page number 146 of Figs. 4 and 6, and the target address is found using the page and the offset 161 of Fig. 6).

13. **With respect to claim 10**, Kikuta et al. in view of Carter et al. disclose the translation lookaside buffer system of claim 2 (see above paragraph 9). Kikuta et al. disclose the limitations wherein the primary lookup logic circuit (MMU 122 of Fig. 2) is further configured to store a new primary tag for the current target address in the primary tag store in the event of a primary cache miss (column 9, lines 47-50, line 66 -- column 10, line 5) thereby designating a location in the primary cluster cache for storing a new cluster corresponding to the new primary tag (column 10, lines 4-5), wherein the new primary tag encodes a minimum-size range of the virtual address space (the virtual page number field, or tag, is composed of a number of bits as in column 9, lines 5-6, and the minimum range of addresses the tag encodes is defined by the number of bits in this field).

14. **With respect to claim 11**, Kikuta et al. in view of Carter et al. disclose the translation lookaside buffer system of claim 10 (see above paragraph 13). Kikuta et al. disclose the limitations wherein upon storing the new cluster in the primary cluster cache, the new primary tag is updated to reflect the range of virtual addresses for which

the new cluster provides a mapping (column 10, lines 4-5; the range is reflected in the number of bits in the virtual page number field).

15. **With respect to claim 12**, Kikuta et al. in view of Carter et al. disclose the translation lookaside buffer system of claim 1 (see above paragraph 8). Kikuta et al. disclose the limitation further comprising a buffer module (load and store queues 127 and 128 of Fig. 2) configured to receive translation requests from the primary lookup logic circuit and to store the received translation requests for subsequent forwarding to the translation module (column 10, lines 7-11).

16. **With respect to claim 22**, Kikuta et al. disclose a method for translating virtual addresses to physical addresses, the method comprising the acts of:

- defining a plurality of clusters (entries in TLB 126 of Figs. 2 and 4; global TLB of Figs. 7 and 9), wherein each cluster provides a mapping to a respective physical address for each virtual address in a range of virtual address space (column 3, lines 36-43; column 8, line 66 -- column 9, line 6);
- receiving a first translation request from a client (Figs. 7 and 9 show plurality of clients as processors 1-3), the first translation request including a first target virtual address (column 9, lines 28-33);
- storing a first tag in a tag store (TLB 126 of Figs. 2 and 4; global TLB of Figs. 7 and 9; column 10, lines 4-5), the first tag encoding a minimum size range of the virtual address space that includes the first target virtual address space (the virtual page number field, or tag, is composed of a number of bits as in column 9, lines 5-6, and the minimum range of addresses the tag encodes is defined by the

number of bits in this field), the first tag further identifying a first location in a primary cluster store (column 8, line 66 -- column 9, line 6, where the tag is the virtual page number, and each location in the TLB is supplied with an entry 140 of Fig. 3);

- associating the first request with the first location in the primary cluster store (column 9, lines 33-38);
- fetching a first cluster that maps a range of the virtual address space that includes the first target virtual address into the first location in the primary cluster store (column 9, line 66 -- column 10, line 4); and
- translating the first target virtual address to a physical address based at least in part on the mapping provided by the first cluster (column 9, lines 33-35).

Kikuta et al. do not expressly disclose the limitation wherein respective sizes of the ranges mapped by different ones of the clusters are different.

However, Carter et al. disclose the limitation wherein the respective ranges of different ones of the clusters have different sizes (column 4, lines 40-46; column 6, lines 32-35).

Kikuta et al. and Carter et al. are analogous art because they are from the same field of endeavor, namely virtual-to-physical address translation in a multiprocessor system.

At the time of invention, it would have been obvious to a person of ordinary skill in the art to combine the TLB system of Kikuta et al. with the variable segment size of Carter et al. The motivation for doing so would have been because "this segmentation

and access control system provides flexibility to the user, while still permitting strictly enforced security" (column 6, lines 42-44).

Therefore, it would have been obvious to a person of ordinary skill in the art to combine Kikuta et al. with Carter et al. for the benefit of a TLB system in a network with variably-sized segments to obtain the invention as specified in claim 22.

17. **With respect to claim 24**, Kikuta et al. in view of Carter et al. disclose the method of claim 22 (see above paragraph 16). Kikuta et al. disclose the limitation further comprising the act of: subsequently to storing the first tag, modifying the first tag to encode the range of the virtual address space mapped by the first cluster (column 9, line 66 -- column 10, line 5; the modification takes place when the new tag is stored in the TLB).

18. **With respect to claim 25**, Kikuta et al. in view of Carter et al. disclose the method of claim 22 (see above paragraph 16). Kikuta et al. disclose the limitations further comprising the acts of:

- subsequently to receiving the first translation request, receiving a second translation request from the client, the second translation request including a second target virtual address (column 9, lines 28-33);
- determining whether the second target virtual address matches the first tag (column 9, lines 33, 47-49); and
- in the event that the second target virtual address matches the first tag:
 - associating the second request with the first location in the primary cluster store (column 9, lines 33-38);

- retrieving the first cluster from the first location in the primary cluster store (column 9, lines 33-38); and
- translating the second target virtual address to a physical address based at least in part on the mapping provided by the first cluster (column 9, lines 33-35)..

19. **Claims 3-4 and 9** are rejected under 35 U.S.C. 103(a) as being unpatentable over Kikuta et al. (US 6,260,131) and Carter et al. (US 6,003,123) as applied to claims 1-2, 5-7, 10-12, 22 and 24-25 above (see paragraphs 8-18), and further in view of Sheets et al. (US 2005/0044340).

20. **With respect to claim 3**, Kikuta et al. in view of Carter et al. disclose the translation lookaside buffer system of claim 2 (see above paragraph 9). Kikuta et al. in view of Carter et al. do not disclose expressly the limitation wherein each of the plurality of clients is allocated a non-overlapping portion of the primary tag store.

However, Sheets et al. disclose the limitation wherein each of the plurality of clients is allocated a non-overlapping portion of the primary tag store (Fig. 6B) [0026, lines 8-12; 0045; 0047, lines 6-11; the RTT is a translation table that stores tags used as indexes into the translation table].

Sheets et al., Kikuta et al. and Carter et al. are analogous art because they are from the same field of endeavor, namely virtual-to-physical address translation in a client system.

At the time of invention, it would have been obvious to a person of ordinary skill in the art to combine the partitioned TLB of Sheets et al. with the TLB of Kikuta et al.

and the variable segment size of Carter et al. The motivation for doing so would have been to allow "the use of standard memory mapping techniques" [0053, lines 8-12].

Therefore, it would have been obvious to a person of ordinary skill in the art to combine Sheets et al. with Kikuta et al. and Carter et al. for the benefit of a TLB system in a client system with a TLB partitioned into separate client blocks to obtain the invention as specified in claim 3.

21. **With respect to claim 4**, Kikuta et al. in view of Carter et al. disclose the translation lookaside buffer system of claim 2 (see above paragraph 9). Kikuta et al. in view of Carter et al. do not disclose expressly the limitation wherein one of the plurality of clients generates translation requests for a plurality of address streams, and wherein the portion of the primary tag store allocated to the one of the clients is subdivided among the address streams.

However, Sheets et al. disclose the limitation wherein one of the plurality of clients generates translation requests for a plurality of address streams [0028, lines 1-2], and wherein the portion of the primary tag store allocated to the one of the clients is subdivided among the address streams [0049].

Sheets et al., Kikuta et al. and Carter et al. are analogous art because they are from the same field of endeavor, namely virtual-to-physical address translation in a client system.

At the time of invention, it would have been obvious to a person of ordinary skill in the art to combine the partitioned TLB of Sheets et al. with the TLB of Kikuta et al.

and the variable segment size of Carter et al. The motivation for doing so would have been to allow "the use of standard memory mapping techniques" [0053, lines 8-12].

Therefore, it would have been obvious to a person of ordinary skill in the art to combine Sheets et al. with Kikuta et al. and Carter et al. for the benefit of a TLB system in a client system with a TLB partitioned into separate client blocks to obtain the invention as specified in claim 4.

22. **With respect to claim 9**, Kikuta et al. in view of Carter et al. disclose the translation lookaside buffer system of claim 2 (see above paragraph 9). Kikuta et al. in view of Carter et al. do not disclose expressly the limitation wherein a non-overlapping portion of the primary tag store is allocated to each client.

However, Sheets et al. disclose the limitation wherein a non-overlapping portion of the primary tag store is allocated to each client (Fig. 6B) [0026, lines 8-12; 0045; 0047, lines 6-11; the RTT is a translation table that stores tags used as indexes into the translation table].

Sheets et al., Kikuta et al. and Carter et al. are analogous art because they are from the same field of endeavor, namely virtual-to-physical address translation in a client system.

At the time of invention, it would have been obvious to a person of ordinary skill in the art to combine the partitioned TLB of Sheets et al. with the TLB of Kikuta et al. and the variable segment size of Carter et al. The motivation for doing so would have been to allow "the use of standard memory mapping techniques" [0053, lines 8-12].

Therefore, it would have been obvious to a person of ordinary skill in the art to combine Sheets et al. with Kikuta et al. and Carter et al. for the benefit of a TLB system in a client system with a TLB partitioned into separate client blocks to obtain the invention as specified in claim 9.

23. **Claims 8 and 23** are rejected under 35 U.S.C. 103(a) as being unpatentable over Kikuta et al. (US 6,260,131) and Carter et al. (US 6,003,123) as applied to claims 1-2, 5-7, 10-12, 22 and 24-25 above (see paragraphs 8-18), and further in view of Bungion et al. (US 6,075,938).

24. **With respect to claim 8**, Kikuta et al. in view of Carter et al. disclose the translation lookaside buffer system of claim 1 (see above paragraph 8). Kikuta et al. in view of Carter et al. do not disclose expressly the limitation wherein the secondary cluster cache further includes: a cluster fetch circuit configured to receive the target virtual address in the event that a matching secondary tag is not identified by the secondary lookup logic circuit and to obtain a cluster corresponding to the target virtual address from a cluster table.

However, Bungion et al. disclose the limitation wherein the secondary cluster cache (second-level TLB "l2tlb" of Fig. 3; column 14, line 7) further includes: a cluster fetch circuit (TLB miss handler, column 14, line 6) configured to receive the target virtual address in the event that a matching secondary tag is not identified by the secondary lookup logic circuit (column 14, lines 5-10) and to obtain a cluster corresponding to the target virtual address from a cluster table (column 14, lines 10-18).

Bungion et al., Kikuta et al. and Carter et al. are analogous art because they are from the same field of endeavor, namely virtual-to-physical address translation.

At the time of invention, it would have been obvious to a person of ordinary skill in the art to combine the lookup of a TLB entry after a second-level TLB miss with the TLB of Kikuta et al. and the variable segment size of Carter et al. The motivation for doing so would have been so "that virtual machines appear to have much larger TLBs than the MIPS processors" (column 13, lines 1-7).

Therefore, it would have been obvious to a person of ordinary skill in the art to combine Bungion et al. with Kikuta et al. and Carter et al. for the benefit of a TLB system in a client system with variably-sized segments and a second-level TLB miss handler to obtain the invention as specified in claim 8.

25. **With respect to claim 23**, Kikuta et al. in view of Carter et al. disclose the method of claim 22 (see above paragraph 16). Kikuta et al. disclose the limitation wherein the act of fetching the first cluster includes the acts of:

- searching for the corresponding cluster in a secondary cluster cache (page table 129 of Fig. 2) based on the target virtual address (column 9, lines 57-59).

Kikuta et al. do not disclose the limitation wherein

- in the event that the corresponding cluster is not found in the secondary cluster cache, fetching the corresponding cluster from a cluster table.

However, Bungion et al. disclose the limitation wherein

- in the event that the corresponding cluster is not found in the secondary cluster cache (second-level TLB "l2tlb" of Fig. 3; column 14, line 7), fetching the corresponding cluster from a cluster table (column 14, lines 5-18).

Bungion et al., Kikuta et al. and Carter et al. are analogous art because they are from the same field of endeavor, namely virtual-to-physical address translation.

At the time of invention, it would have been obvious to a person of ordinary skill in the art to combine the lookup of a TLB entry after a remote TLB miss with the TLB of Kikuta et al. and the variable segment size of Carter et al. The motivation for doing so would have been so "that virtual machines appear to have much larger TLBs than the MIPS processors" (column 13, lines 1-7).

Therefore, it would have been obvious to a person of ordinary skill in the art to combine Bungion et al. with Kikuta et al. and Carter et al. for the benefit of a TLB system in a client system with variably-sized segments and a second-level TLB miss handler to obtain the invention as specified in claim 23.

26. **Claim 13** is rejected under 35 U.S.C. 103(a) as being unpatentable over Kikuta et al. (US 6,260,131) and Carter et al. (US 6,003,123) as applied to claims 1-2, 5-7 and 11-12 above (see paragraphs 8-15), and further in view of Heaslip et al. (US 5,930,832).

27. **With respect to claim 13**, Kikuta et al. in view of Carter et al. disclose the translation lookaside buffer system of claim 1 (see above paragraph 8). Kikuta et al. in view of Carter et al. do not disclose expressly the limitations wherein the primary lookup logic circuit is further configured to process a selected translation request before the

cluster for a previously selected translation request is present in the primary cluster store location associated with the previously selected translation request.

However, Heaslip et al. disclose the limitations wherein the primary lookup logic circuit is further configured to process a selected translation request before the cluster for a previously selected translation request is present in the primary cluster store location associated with the previously selected translation request (column 5, lines 38-45; column 7, lines 8-24).

Kikuta et al., Carter et al. and Heaslip et al. are analogous art because they are from the same field of endeavor, namely translation lookaside buffers.

At the time of invention, it would have been obvious to a person of ordinary skill in the art to combine the storage of a missed load instruction for later processing of Heaslip et al. with the TLB of Kikuta et al. and Carter et al.

The motivation for doing so would have been to avoid "the need for such complex and space consuming logic for recovering from a SLB/TLB miss" (column 6, lines 22-24).

Therefore, it would have been obvious to a person of ordinary skill in the art to combine Heaslip et al. with Kikuta et al. and Carter et al. for the benefit of a TLB with storage of a missed load instruction for later processing to obtain the invention as specified in claim 13.

28. **Claims 14-19** are rejected under 35 U.S.C. 103(a) as being unpatentable over Kikuta et al. (US 6,260,131) and Carter et al. (US 6,003,123) as applied to claims 1-2,

5-7 and 11-12 above (see paragraphs 8-15), and further in view of Nielsen et al. (US 6,104,417).

29. **With respect to claim 14**, Kikuta et al. in view of Carter et al. disclose the translation lookaside buffer system of claim 12 (see above paragraph 15). Kikuta et al. in view of Carter et al. do not disclose expressly the limitation wherein the buffer module includes a separate request queue for each of the plurality of clients.

However, Nielsen et al. disclose the limitation wherein the buffer module includes a separate request queue for each of the plurality of clients (client queues of Fig. 4, top drawing; column 10, lines 59-64).

Nielsen et al., Kikuta et al. and Carter et al. are analogous art because they are from the same field of endeavor, namely virtual-to-physical address translation in a client system.

At the time of invention, it would have been obvious to a person of ordinary skill in the art to combine the client queues of Nielsen et al. with the TLB of Kikuta et al. and the variable segment size of Carter et al. The motivation for doing so would have been to ensure a specific ratio of requests from particular clients (column 11, lines 52-57).

Therefore, it would have been obvious to a person of ordinary skill in the art to combine Nielsen et al. with Kikuta et al. and Carter et al. for the benefit of a TLB system in a client system with variably-sized segments and client queues to obtain the invention as specified in claim 14.

30. **With respect to claim 15**, Kikuta et al. and Carter et al. in view of Nielsen et al. disclose the translation lookaside buffer system of claim 14 (see above paragraph 29).

Kikuta et al. and Carter et al. do not disclose expressly the limitation wherein the request queues are configured such that a sequential order of receipt among translation requests from a same one of the clients is preserved and a sequential order of receipt among translation requests from different ones of the clients is not preserved.

However, Nielsen et al. disclose the limitation wherein the request queues are configured such that a sequential order of receipt among translation requests from a same one of the clients is preserved (column 10, lines 14-17) and a sequential order of receipt among translation requests from different ones of the clients is not preserved (column 10, lines 17-20).

Nielsen et al., Kikuta et al. and Carter et al. are analogous art because they are from the same field of endeavor, namely virtual-to-physical address translation in a client system.

At the time of invention, it would have been obvious to a person of ordinary skill in the art to combine the client queues of Nielsen et al. with the TLB of Kikuta et al. and the variable segment size of Carter et al. The motivation for doing so would have been to ensure a specific ratio of requests from particular clients (column 11, lines 52-57).

Therefore, it would have been obvious to a person of ordinary skill in the art to combine Nielsen et al. with Kikuta et al. and Carter et al. for the benefit of a TLB system in a client system with variably-sized segments and client queues to obtain the invention as specified in claim 15.

31. **With respect to claim 16**, Kikuta et al. in view of Carter et al. disclose the translation lookaside buffer system of claim 12 (see above paragraph 15) and the limitation further comprising:

- a ready logic circuit (MMU 122 of Fig. 2) configured to identify as a ready request at least one translation request stored in the buffer module for which a cluster is present in the associated location in the primary cluster cache (column 9, lines 40-46).

Kikuta et al. in view of Carter et al. do not disclose expressly the limitation further comprising:

- a selection circuit configured to select one of the ready requests and to provide the selected ready request to the translation module.

However, Nielsen et al. disclose the limitation further comprising:

- a selection circuit (arbitration logic, column 10, line 17) configured to select one of the ready requests and to provide the selected ready request to the translation module (column 10, lines 17-20).

Nielsen et al., Kikuta et al. and Carter et al. are analogous art because they are from the same field of endeavor, namely virtual-to-physical address translation in a client system.

At the time of invention, it would have been obvious to a person of ordinary skill in the art to combine an arbitration circuit that chooses one of the translation requests of Nielsen et al. with the TLB of Kikuta et al. and the variable segment size of Carter et al.

The motivation for doing so would have been to ensure a specific ratio of requests from particular clients (column 11, lines 52-57).

Therefore, it would have been obvious to a person of ordinary skill in the art to combine Nielsen et al. with Kikuta et al. and Carter et al. for the benefit of a TLB system in a client system with variably-sized segments and arbitration logic that chooses one of the translation requests to obtain the invention as specified in claim 16.

32. **With respect to claim 17**, Kikuta et al. in view of Carter et al. disclose the translation lookaside buffer system of claim 1 (see above paragraph 8). Kikuta et al. in view of Carter et al. do not disclose expressly the limitation wherein the input arbitration module is further configured to select a translation request based at least in part on availability of downstream resources needed to process each translation request.

However, Nielsen et al. disclose the limitation wherein the input arbitration module is further configured to select a translation request based at least in part on availability of downstream resources needed to process each translation request (column 10, lines 16-17, lines 56-58, line 64 -- column 11, line 9).

Nielsen et al., Kikuta et al. and Carter et al. are analogous art because they are from the same field of endeavor, namely virtual-to-physical address translation in a client system.

At the time of invention, it would have been obvious to a person of ordinary skill in the art to combine checking for the availability of downstream resources of Nielsen et al. with the TLB of Kikuta et al. and the variable segment size of Carter et al. The

motivation for doing so would have been to ensure a specific ratio of requests from particular clients (column 11, lines 52-57).

Therefore, it would have been obvious to a person of ordinary skill in the art to combine Nielsen et al. with Kikuta et al. and Carter et al. for the benefit of a TLB system in a client system with variably-sized segments and checking for the availability of downstream resources to obtain the invention as specified in claim 17.

33. **With respect to claim 18**, Kikuta et al. and Carter et al. in view of Nielsen et al. disclose the translation lookaside buffer system of claim 17 (see above paragraph 32). Kikuta et al. and Carter et al. do not disclose expressly the limitation wherein the needed downstream resources include a storage location in the primary cluster store.

However, Nielsen et al. disclose the limitation wherein the needed downstream resources include a storage location in the primary cluster store (column 10, lines 56-58; column 11, lines 4-9; the primary cluster store is the decode logic).

Nielsen et al., Kikuta et al. and Carter et al. are analogous art because they are from the same field of endeavor, namely virtual-to-physical address translation in a client system.

At the time of invention, it would have been obvious to a person of ordinary skill in the art to combine checking for the availability of downstream resources of Nielsen et al. with the TLB of Kikuta et al. and the variable segment size of Carter et al. The motivation for doing so would have been to ensure a specific ratio of requests from particular clients (column 11, lines 52-57).

Therefore, it would have been obvious to a person of ordinary skill in the art to combine Nielsen et al. with Kikuta et al. and Carter et al. for the benefit of a TLB system in a client system with variably-sized segments and checking for the availability of downstream resources to obtain the invention as specified in claim 18.

34. **With respect to claim 19**, Kikuta et al. and Carter et al. in view of Nielsen et al. disclose the translation lookaside buffer system of claim 17 (see above paragraph 32). Kikuta et al. in view of Carter et al. do not disclose expressly the limitation wherein the input arbitration module is further configured to determine whether the needed downstream resources are available or not available based at least in part on which of the clients generated the proposed request.

However, Nielsen et al. disclose the limitation wherein the input arbitration module is further configured to determine whether the needed downstream resources are available or not available based at least in part on which of the clients generated the proposed request (column 10, lines 16-17, lines 64-66).

Nielsen et al., Kikuta et al. and Carter et al. are analogous art because they are from the same field of endeavor, namely virtual-to-physical address translation in a client system.

At the time of invention, it would have been obvious to a person of ordinary skill in the art to combine checking for the availability of downstream resources of Nielsen et al. with the TLB of Kikuta et al. and the variable segment size of Carter et al. The motivation for doing so would have been to ensure a specific ratio of requests from particular clients (column 11, lines 52-57).

Therefore, it would have been obvious to a person of ordinary skill in the art to combine Nielsen et al. with Kikuta et al. and Carter et al. for the benefit of a TLB system in a client system with variably-sized segments and checking for the availability of downstream resources to obtain the invention as specified in claim 19.

35. **Claims 20-21 and 26** are rejected under 35 U.S.C. 103(a) as being unpatentable over Kikuta et al. (US 6,260,131) and Carter et al. (US 6,003,123) as applied to claims 1-2, 5-7 and 11-12 above (see paragraphs 8-15), and further in view of Vishin et al. (US 5,860,146).

36. **With respect to claim 20**, Kikuta et al. in view of Carter et al. disclose the translation lookaside buffer system of claim 12 (see above paragraph 15). Kikuta et al. in view of Carter et al. do not disclose expressly the limitations wherein the translation requests received from the clients include a first request from a first one of the clients, the first request having a transparent translation mode, and wherein the system further comprises: a memory interface circuit configured to transmit a memory access request corresponding to the first request to a system memory device, the memory access request including the physical address provided by the translation module in response to the first request.

However, Vishin et al. disclose the limitations wherein the translation requests received from the clients (processors 104 of Fig. 1) include a first request from a first one of the clients (column 2, lines 2-4), the first request having a transparent translation mode (column 1, lines 13-18), and wherein the system further comprises: a memory interface circuit (bus 106 of Fig. 1; column 1, lines 15-16) configured to transmit a

memory access request corresponding to the first request to a system memory device (column 2, lines 20-27), the memory access request including the physical address provided by the translation module in response to the first request (column 2, lines 20-27).

Vishin et al., Kikuta et al. and Carter et al. are analogous art because they are from the same field of endeavor, namely virtual-to-physical address translation in a client system.

At the time of invention, it would have been obvious to a person of ordinary skill in the art to combine the local translation mode of Vishin et al. with the TLB of Kikuta et al. and the variable segment size of Carter et al. The motivation for doing so would have been because a TLB is "like a conventional cache memory" (column 1, lines 64-67), which is faster than normal memory.

Therefore, it would have been obvious to a person of ordinary skill in the art to combine Vishin et al. with Kikuta et al. and Carter et al. for the benefit of a TLB system in a client system with variably-sized segments and a local translation mode to obtain the invention as specified in claim 20.

37. **With respect to claim 21**, Kikuta et al. in view of Carter et al. disclose the translation lookaside buffer system of claim 12 (see above paragraph 15). Kikuta et al. in view of Carter et al. do not disclose expressly the limitations wherein the translation requests received from the clients further include a second request from a second one of the clients, the second request having a visible translation mode, and wherein the

memory interface circuit is further configured to transmit to the second client the physical address provided by the translation module in response to the second request.

However, Vishin et al. disclose the limitations wherein the translation requests received from the clients further include a second request from a second one of the clients (column 2, lines 2-4; column 4, lines 59-62), the second request having a visible translation mode (column 3, lines 17-30), and wherein the memory interface circuit is further configured to transmit to the second client the physical address provided by the translation module in response to the second request (column 3, lines 26-30; the second client is another cluster 102 on the network 114 of Fig. 9).

Vishin et al., Kikuta et al. and Carter et al. are analogous art because they are from the same field of endeavor, namely virtual-to-physical address translation in a client system.

At the time of invention, it would have been obvious to a person of ordinary skill in the art to combine the remote translation mode of Vishin et al. with the TLB of Kikuta et al. and the variable segment size of Carter et al. The motivation for doing so would have been because the hardware remote translation mode of Vishin et al. "speeds up the process of converting a physical address into a remote physical address considerably" (column 4, lines 55-58).

Therefore, it would have been obvious to a person of ordinary skill in the art to combine Vishin et al. with Kikuta et al. and Carter et al. for the benefit of a TLB system in a client system with variably-sized segments and a remote translation mode to obtain the invention as specified in claim 21.

38. **With respect to claim 26**, Kikuta et al. in view of Carter et al. disclose the translation lookaside buffer system of claim 12 (see above paragraph 15). Kikuta et al. in view of Carter et al. do not disclose expressly the limitations wherein

- determining whether the first translation request is in a transparent translation mode or a visible translation mode;
- in the event that the first translation request is in the transparent translation mode, transmitting a corresponding memory access operation to a system memory device, wherein the corresponding memory access operation includes the translated physical address; and
- in the event that the first translation request is in the visible translation mode, transmitting the translated physical address to the client.

However, Vishin et al. disclose the limitations wherein

- determining whether the first translation request is in a transparent translation mode or a visible translation mode (column 4, lines 59-62);
- in the event that the first translation request is in the transparent translation mode, transmitting a corresponding memory access operation to a system memory device (column 2, lines 20-27), wherein the corresponding memory access operation includes the translated physical address (column 2, lines 20-27); and
- in the event that the first translation request is in the visible translation mode, transmitting the translated physical address to the client (column 3, lines 26-30; the second client is another cluster 102 on the network 114 of Fig. 9).

Vishin et al., Kikuta et al. and Carter et al. are analogous art because they are from the same field of endeavor, namely virtual-to-physical address translation in a client system.

At the time of invention, it would have been obvious to a person of ordinary skill in the art to combine the local and remote translation modes of Vishin et al. with the TLB of Kikuta et al. and the variable segment size of Carter et al. The motivation for doing so would have been because the hardware remote translation mode of Vishin et al. "speeds up the process of converting a physical address into a remote physical address considerably" (column 4, lines 55-58).

Therefore, it would have been obvious to a person of ordinary skill in the art to combine Vishin et al. with Kikuta et al. and Carter et al. for the benefit of a TLB system in a client system with variably-sized segments and local and remote translation modes to obtain the invention as specified in claim 26.

Conclusion

39. The prior art made of record and not relied upon is considered pertinent to applicant's disclosure.

- Hass et al. (US 2005/0055510) teach a shared TLB in a multiprocessor system.
- McGrath (US 6,880,068) teaches a TLB variably-sized pages.
- Gaertner et al. (US 6,766,434) teach a multi-level shared TLB in a multiprocessor system.
- Willis et al. (US 6,728,858) also teach a shared TLB in a multiprocessor system.

Art Unit: 2187

- Kessler et al. (US 6,715,057) also teach a TLB with variably-sized pages.
- Nijhawan et al. (US 6,374,341) also teach a TLB with variably-sized pages.

40. Any inquiry concerning this communication or earlier communications from the examiner should be directed to James Golden whose telephone number is 571-272-5628. The examiner can normally be reached on Monday-Friday, 8:30 AM - 5:30 PM.


If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Donald Sparks can be reached on 571-272-4201. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

James R. Golden
Patent Examiner
Art Unit 2187



April 10, 2006



Brian R. Peg
Primary Examiner
Art 2187